

ABSTRACT OF THE DISCLOSURE

[0050] A method, apparatus, and computer readable medium for processing a data stream is described. In one example, a set of elements of a data stream are received. The set of elements are stored in a memory as a hierarchy of nodes. Each of the nodes includes frequency data associated with either an element in the set of elements or a prefix of an element in the set of elements. A set of hierarchical heavy hitters is then identified among the nodes in the hierarchy. The frequency data of each of the hierarchical heavy hitter nodes, after discounting any portion thereof attributed to a descendent hierarchical heavy hitter node in said set of hierarchical heavy hitter nodes, being greater than or equal to a fraction of the number of elements in the set of elements.